

<https://itorizon.com/job/senior-agentic-ai-platform-engineer-mcp-agent-runtime/>

## Senior Agentic AI Platform Engineer – MCP & Agent Runtime

### Description About Us

ITOrizon is a global consulting and technology company that helps enterprises design, implement, and optimize complex supply chain and digital transformation initiatives. Headquartered in Atlanta, USA, with offices in India (Bengaluru) and the UAE (Sharjah), we partner with global clients across retail, manufacturing, logistics, and distribution sectors.

We combine deep domain expertise with modern technology to deliver practical, scalable solutions. Our teams work across strategy, implementation, and managed services — helping organizations adopt leading platforms such as Oracle, Manhattan, Blue Yonder, and our next-generation composable enterprise platform, Karolium.

### Role Overview

We are looking for a Senior Agentic AI Platform Engineer to build the core platform infrastructure for AI agents, including runtimes, SDKs, MCP integrations, and evaluation frameworks.

This is a coding-first, platform-focused role—you will not build isolated AI use cases, but instead create reusable systems, developer tooling, and scalable infrastructure that enable teams to build and operate AI agents reliably across the enterprise.

### Key Responsibilities

- Design, develop, and deploy machine learning and statistical models for enterprise use cases
- Perform data exploration, feature engineering, and model evaluation on structured and unstructured data
- Build predictive, prescriptive, and descriptive analytics solutions
- Work with large-scale datasets, real-time data pipelines, and cloud data platforms
- Collaborate with engineering teams to productionize models using MLOps practices
- Develop and optimize models across:
  - Classical ML (regression, classification, clustering)
  - Time-series forecasting & anomaly detection
  - Deep learning & NLP
- Ensure model performance, scalability, explainability, and reliability
- Implement model monitoring, versioning, and lifecycle management
- Contribute to AI architecture, experimentation frameworks, and innovation initiatives
- Stay updated with AI/ML and GenAI advancements and apply them to real-world problems

**Hiring organization**  
ITOrizon

**Employment Type**  
Full-time

**Experience**  
5+ Years

**Role Category**  
Agentic AI Engineer

**Role**  
Senior

**Job Location**  
Bangalore, India, India

**Date posted**  
April 9, 2026

## **Agent Runtime & Platform Foundations**

- Build and maintain a scalable Agent Runtime with:
  - Orchestration primitives
  - Tool routing, retries, fallbacks
  - Memory and state management
  - Guardrails and safety controls
- Design reusable agent execution frameworks for multi-tenant environments

## **Developer Platform & SDKs**

- Develop a developer-first Agent SDK including:
  - Libraries, templates, and reference implementations
- Standardize how internal teams build, deploy, and manage AI agents

## **MCP (Model Context Protocol) & Integrations**

- Implement MCP exposure layers for secure tool access
- Build MCP servers and connectors for:
  - Enterprise systems
  - Internal platform services
- Ensure authentication, authorization, and auditability

## **RAG & Knowledge Systems**

- Design and implement platform-grade RAG pipelines, including:
  - Data ingestion and chunking strategies
  - Embedding pipelines and vector databases
  - Retrieval, reranking, and grounding mechanisms

## **Evaluation, Observability & Reliability**

- Build evaluation frameworks:
  - Offline evaluation
  - Prompt/tool regression testing
  - Quality metrics and drift detection
- Implement observability systems:
  - Tracing, logging, telemetry
  - Cost/token optimization
  - Latency and error monitoring

## **Enterprise Platform Readiness**

- Ensure multi-tenancy and platform security, including:
  - Tenant isolation
  - Environment separation (Dev/QA/Prod)
  - Secret management
  - Role-based access control (RBAC)
- Integrate with enterprise systems:
  - APIs, integration frameworks, security layers

## **Engineering Excellence**

- Drive clean code practices, design patterns, and architecture standards
- Conduct code reviews and technical documentation
- Contribute to platform roadmap and innovation initiatives

## **Required Skills & Expertise**

## **Core Engineering (Coding-First)**

- Strong hands-on expertise in Python (primary language)
- Experience with Java / Node.js / Go (multi-language ecosystem)
- Strong fundamentals in:
  - API design
  - Clean architecture & design patterns
  - Performance optimization
- Experience building:
  - Microservices
  - Secure APIs
  - Integration platforms

## **Agentic AI & LLM Engineering**

- Experience building tool-using AI agents, including:
  - Function/tool calling
  - Prompt engineering & templates
  - Conversation state & memory
  - Guardrails and structured outputs
- Understanding of MCP concepts and implementations

## **RAG & Data Systems**

- Hands-on experience with:
  - Vector databases & embeddings
  - Retrieval strategies and ranking
  - Grounding and citation technique

## **Cloud & Platform Engineering**

- Experience with AWS and/or Azure
- Familiarity with:
  - Docker & Kubernetes
  - CI/CD pipelines
  - Enterprise security and networking

## **Preferred Certifications**

- AWS Certified Machine Learning – Specialty
- Microsoft Certified: Azure AI Engineer / Data Scientist
- Kubernetes (CKA / CKAD)
- Security certifications (CISSP, CCSP)

## **Good-to-Have Skills**

- Experience building internal developer platforms (IDPs) or SDKs
- Familiarity with:
  - Multi-agent systems
  - Orchestrator/supervisor patterns
- Knowledge of Responsible AI, governance, and audit frameworks

## **Ideal Candidate Profile**

- 5+ years of strong engineering experience with production system ownership
- Deep expertise in Python and AI platform development
- Experience building scalable, reusable platform components

- Strong platform mindset (enabling teams, not just solving use cases)
- High adaptability in rapidly evolving AI ecosystems

### **Why Join Us?**

- Build core AI agent platform infrastructure (runtime, SDKs, MCP)
- Work on enterprise-scale AI challenges:
  - Multi-tenancy
  - Security
  - Observability
  - Governance
- Be part of a cutting-edge AI-first platform roadmap

### **How to Apply:**

Email your Resume to [engage@itorizon.com](mailto:engage@itorizon.com)